Machine Protection and Operation for LHC

J. Wenninger
CERN, Geneva, Switzerland

Abstract
Since 2010 the Large Hadron Collider (LHC) is the accelerator with the highest stored energy per beam, with a record of 140 MJ at a beam energy of 4 TeV, almost a factor of 50 higher than other accelerators. With such a high stored energy, machine protection aspects set the boundary conditions for operation during all phases of the machine cycle. Only the low-intensity commissioning beams can be considered as relatively safe. This document discusses the interplay of machine operation and machine protection at the LHC, from commissioning to regular operation.

Keywords
Operation; machine protection; beam loss; LHC.

arXiv:1608.03113v1 [physics.acc-ph] 10 Aug 2016

1 Introduction

The Large Hadron Collider (LHC) is the last in the series of hadron colliders after the ISR (Intersecting Storage Ring), SPS (Super Proton Synchrotron), Tevatron, HERA and RHIC (Relativistic Heavy Ion Collider). The machine elements are installed on average 100 m below the surface in the 26.7 km long accelerator tunnel that housed the Large Electron Positron collider (LEP) between 1989 and 2000 [1, 2]. The ring consists of eight arcs and of eight long straight sections (LSSs). The large particle physics experiments ALICE, ATLAS, CMS and LHCb are installed at interaction points (IPs) in the middle of four LSSs, while the other LSSs house the collimation (or beam cleaning) system, the radio-frequency (RF) system, the beam instrumentation and the beam dumping system. The layout of the LHC is shown in Fig. 1.

A dipole field of 8.3 T is required to bend hadrons with a momentum of 7 TeV/c per unit charge in the tunnel; this is 60% higher than in previous accelerators. Such a magnetic field strength is achieved with superconducting dipole magnets made of NbTi. With a 2-in-1 magnet design the two rings fit inside the 3.8 m diameter LEP tunnel; see Fig. 2. Both rings are accommodated in a single cryostat and the distance between the two vacuum chambers is only 19 cm. The two proton or ion beams circulate in opposite directions in two independent vacuum chambers. Each dipole magnet is 14.3 m long; the associated cryostat is 15 m long. Besides the 1232 dipole magnets that constitute around 85% of each arc, the magnet lattice also includes quadrupole magnets that focus the beam, sextupole magnets to correct chromatic effects and octupoles to stabilize the beam. A total of 8000 superconducting magnets are used to control the two beams.

Eight continuous cryostats with a length of 2.7 km each cool the superconducting magnets to their operating temperature of 1.9 K. After cool down the LHC cryostats contain 130 tons of liquid helium and around 37 000 tons of material are cooled to that temperature. The magnets and the cooling system based on superfluid helium form by far the longest refrigerators on Earth. The cryogenic system has to be extremely reliable; in 2012 the system achieved an overall uptime of 95%.

Around 1600 power converters provide current to the magnets; for the main circuits the peak currents reach 13 kA. The magnetic energy stored in each arc cryostat is around 1 GJ. This energy has to be safely extracted in case one of the magnets quenches, i.e. performs a transition from the superconducting to the normal-conducting state [3]. Large dump resistors that are capable of absorbing the energy are automatically switched into the main circuits in case of a quench.

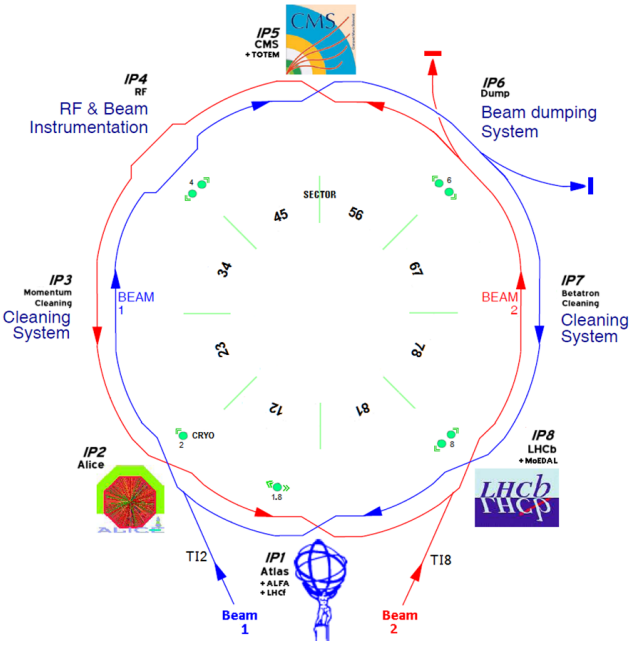

Fig. 1: Layout of the LHC with the eight interaction points labelled IP1 to IP8. The experiments ATLAS, ALICE, CMS and LHCb are installed in IP1, IP2, IP5 and IP8, respectively. Beam 1 is injected close to IP2 and circulates clockwise, Beam 2 is injected close to IP8 and circulates counter-clockwise. The two beams exchange position between outside and inside of the ring at every experiment to ensure that the path length is the same for both beams. The two beam dumps are located around IP6.

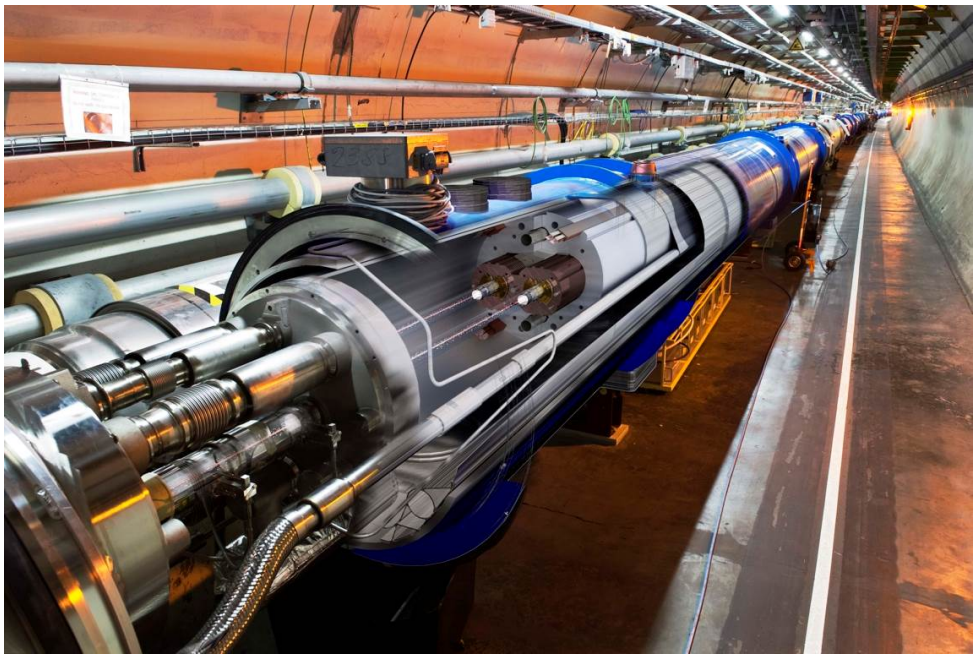

Fig. 2: View of the LHC tunnel with an artistic cut through a LHC dipole magnet, highlighting the twin magnet coils as well as the two vacuum chambers.

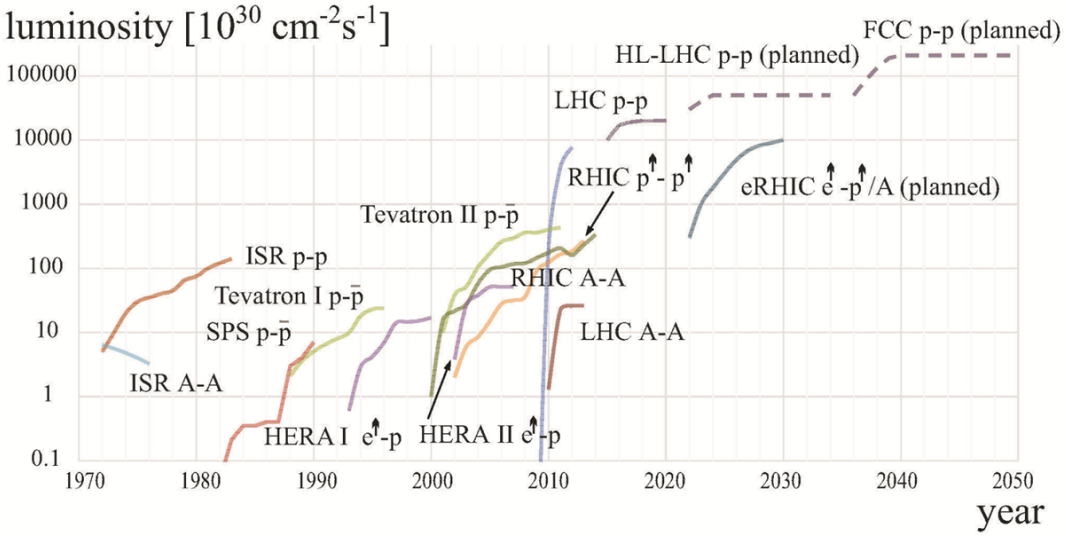

Fig. 3: Evolution of the peak luminosity of hadron colliders since 1970 (image courtesy of W. Fischer, BNL). Future projects like the high-luminosity LHC (HL-LHC) [5, 6], eRHIC and a possible 100 km circumference collider (FCC) [7] are also indicated.

The performance of a collider is characterized by its luminosity $\mathcal{L}$. The event rate of a physical process with cross-section σ (with unit of an area [m^2]) is given by $\sigma \times \mathcal{L}$. The luminosity may be expressed as

$$\mathcal{L} = \frac{kfN^2}{4\pi\sigma_x\sigma_y}, \tag{1}$$

where f is the revolution frequency (11.24 kHz for the LHC), k is the number of bunches, N is the number of particles per bunch and σ_x and σ_y are the horizontal and vertical beam sizes at the collision point. The highest luminosity is achieved with the smallest possible beam cross-section, a large number of bunches and a high bunch population. Up to 2808 bunches can be filled and accelerated in each LHC beam; the minimum distance between bunches is 25 ns or 7.5 m. Each proton bunch consists of up to 1.7×10^{11} protons. The bunches have a typical length of 7 to 10 cm. At the interaction point the transverse rms sizes of the colliding beams are around 20 μm [4]. Figure 3 presents the evolution of hadron collider luminosity over time. The LHC pushes the energy frontier by a factor of seven and the luminosity frontier by a factor of 25; the luminosity gain is mainly obtained with very high beam intensities.

The LHC dipole magnets were produced by three industrial firms and the last dipole magnet was delivered to CERN in November 2006. Each magnet was trained on CERN test benches to a magnetic field of 8.7 T, approximately 5% above the design target. A few training quenches were typically required to reach the nominal field of 8.3 T. Training quenches are due to the release of extremely small amounts of frictional energy (10–100 nJ) due to coil movements when the magnetic field is increased. In June 2007 the first arc (sector) of the LHC was cooled down and ready for commissioning and in April 2008 the last dipole magnet was lowered into the LHC tunnel. One of the essential components of the commissioning phase was the testing of the LHC superconducting magnets and the associated powering and protection equipment. In early 2008 it became apparent that the LHC dipole magnets had to be re-trained to their nominal field; the first magnet quenches appeared at fields corresponding to beam energies of around 5.5 TeV. A training campaign on one arc revealed that the number of required re-training quenches increased rapidly with the magnetic field. The estimated number of quenches required to reach 6.5 TeV is around 140, confirmed during the re-commissioning in 2015, while for 7 TeV the expected number

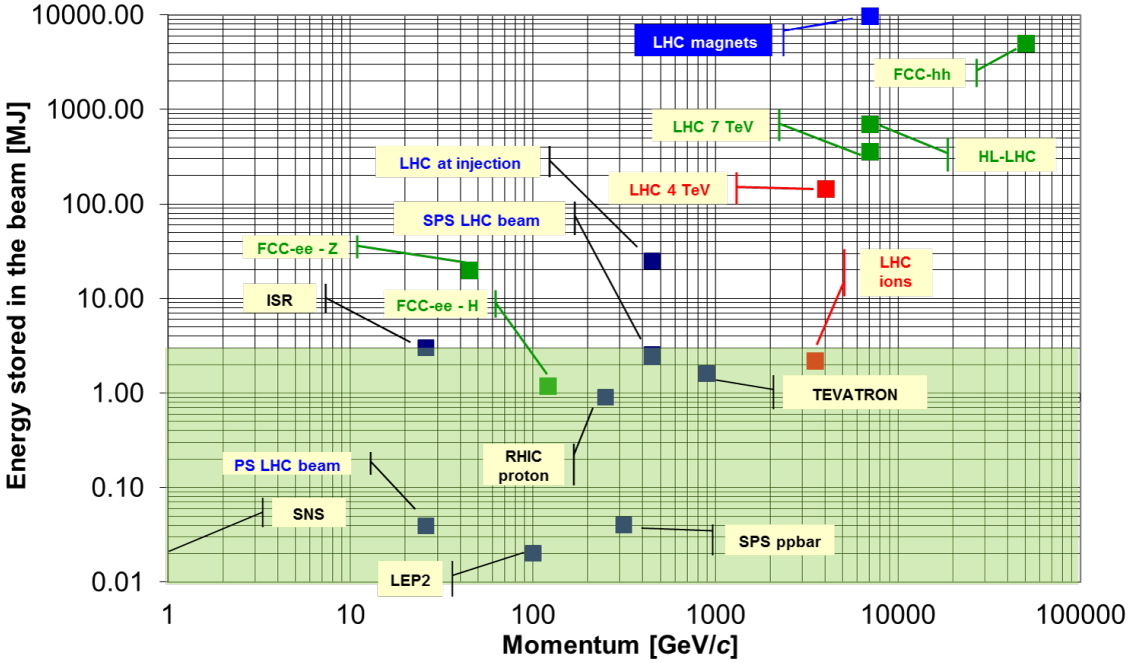

Fig. 4: Stored energy of the beams as a function of the momentum for various accelerators. The area shaded in light green corresponds to the state of the art before the startup of the LHC. All accelerators operate with hadrons except the lepton colliders LEP2 and FCC-ee (the CERN 100 km collider study).

can be as high as 1000. Since such a training campaign would have required a long time, it was decided to lower the energy for the commissioning and first operation phase to 5 TeV [8].

On 10 September 2008 beams were circulating for the first time in both LHC rings. The startup was however brought to an abrupt halt on 19 September 2008 when a defective high-current soldering between two magnets triggered an electric arc that released around 600 MJ of stored magnetic energy [9]. The accelerator was damaged over a distance of around 700 m; 53 magnets had to be replaced or repaired. The beam vacuum was polluted with dust and soot over 2 km. The repair, improvement and re-commissioning of the LHC lasted until November 2009. More details on the incident that did not involve beams are presented in the Appendix. During the repair a systematic soldering problem affecting roughly 15% of all the high-current (13 kA) cable joints was discovered. Since this problem could not be solved immediately, the beam energy had to be limited to 3.5 TeV until a complete repair campaign could be performed. The LHC therefore operated in 2010 and 2011 at this energy, before the energy was increased to 4 TeV in 2012 and 2013 because no magnet was quenched at 3.5 TeV during beam operation, thus lowering the risk of problems with the cable joints. The energy will be pushed to 6.5 TeV and above in 2015 after the repair campaign of 2013 and 2014 [10, 11].

2 Machine protection at the LHC

At 7 TeV each nominal LHC beam stores an energy of 360 MJ. This corresponds to the energy content of 80 kg of explosives and is a hundred times higher than previously achieved accelerator records. The energy stored in the beams of various accelerators is shown in Fig. 4. Machine protection systems (MPSs) with unprecedented safety levels are therefore required to operate the LHC [12,13]. This is achieved with a combination of active protection by equipment and beam parameter monitoring, as well as with passive protection by a large number of collimators as indicated in Fig. 5.

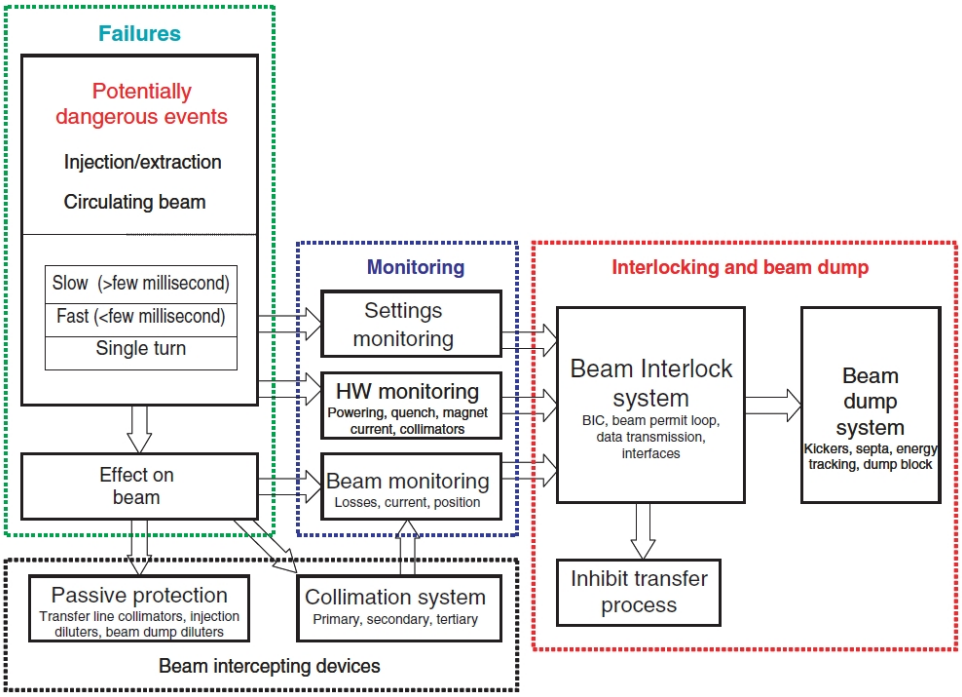

Fig. 5: Schematic outline of the main functions of the LHC MPS. Protection against failures is achieved by active monitoring followed by a dump action or by passive protection with absorbers.

The function of the LHC MPS is to protect the LHC accelerator, its injection and extraction transfer lines, as well as its experiments [14] against a large spectrum of failures [15]. Failures may be detected by monitoring equipment or beam parameters, for example power converter currents and state or beam loss rates. In parallel their impact is mitigated by passive absorbers in the form of collimators and absorbers. In the case where abnormal conditions are detected, interlock signals are sent to the the beam interlock system (BIS) [12, 13] that will trigger a beam dump or stop injection of the beams. A summary of the time-scale of failures and of the reaction time of MPS components is shown in Fig. 6.

The LHC BIS interfaces the sources of interlocks (clients) and the beam dumping and injection kicker systems. There are currently 189 inputs from client systems. The list of connected systems includes:

- the powering interlocks for normal-conducting and superconducting circuits;
- the fast current change detection of electrical circuits with fast failure time constants;
- the electrical circuits of the main experiment magnets;
- the interlock signals from the experiments (excessive rates);
- approximately 4000 beam loss monitors (BLMs) [18];
- the beam position monitors (BPMs);
- the beam screens;
- the collimators and absorbers (positions and temperatures);
- the movable experimental detectors;
- the RF system;
- the beam dumping system (state and trigger unit);
- the injection and aperture kicker systems;
- the personnel access system and its associated beam stoppers;

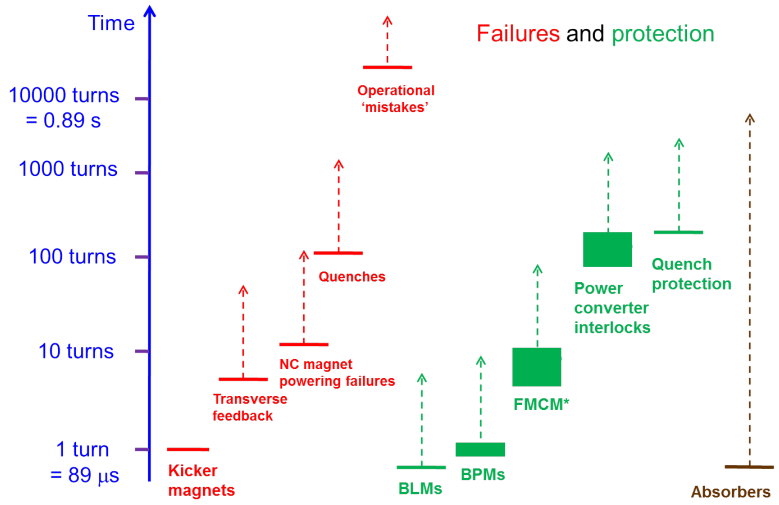

Fig. 6: Time-scales for failures (red) and for the reaction of the active protection systems (green) at the LHC. The symbol FMCM stands for fast magnet current change monitor, a device to detect fast current changes in normal-conduction magnet circuits with very short time constants [16]. The symbol NC stands for normal-conducting magnet.

– the operator inhibit buttons;
– the vacuum valves;
– the programmable beam dump.

Beam interlock controller (BIC) modules are installed in each LHC access point and provide the interface between the client systems and the BIS; see Fig. 7. The BIC units are connected to each other and to the beam dumping system though permit loops. There are two permit loops per beam; the signals in the loops propagate clockwise in one loop and counter-clockwise in the other loop. This arrangement ensures always the fastest possible transmission of a dump request to the dumping system. At the LHC the dump delay can reach around three turns or 270 μs; this delay includes propagation of an interlock on the permit loops, synchronization of the beam dump with the beam abort gap (see below) and one turn to extract the entire beam.

The permit loops are connected directly to the LHC injection interlock system, which is made from the same building blocks (BIC modules). The LHC injection interlocks act on the LHC injection kicker and on the SPS extraction interlock system. This ensures that no beam may be extracted from the SPS injector when the beam permit loops are not armed [17].

The superconducting magnets in the LHC lose their superconducting state (quench) when they are heated with about 10 mJ per cubic centimetre. This is to be compared with the stored energy of many 100 MJ. To prevent unavoidable particle losses from the beam hitting the vacuum chamber within the magnets, possibly leading to sufficiently large energy deposition to trigger a quench [19], a collimation system [20] must intercept the beam losses with very high efficiency. Contrary to previous hadron colliders that used collimators only for experimental background conditions, the LHC cannot operate without its collimation system. Approximately 100 LHC collimators are installed in two long straight sections as shown in Fig. 1. To be able to absorb the energy of the 7 TeV hadrons, the LHC requires a multistage collimation system that intercepts the particles in a four-step process as outlined in Fig. 8, covering the entire six-dimensional phase space, catching particles with large transverse oscillation amplitudes and energy errors [20]. Each collimator is composed of two parallel, typically 1.2 m long material blocks (jaws) that define a variable gap through which the beams circulate. The smallest gap widths are roughly 2 mm.

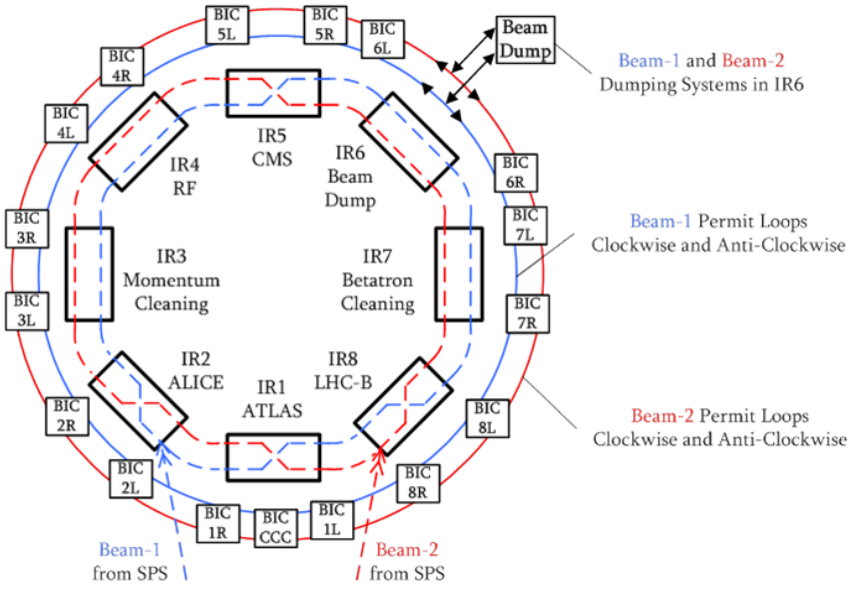

Fig. 7: Schematic layout of the LHC beam interlock system with 17 beam interlock controllers (BICs) distributed around the ring (image courtesy of B. Todd, CERN). The BIC modules are connected to beam permit loops that interface with the LHC beam dumping system.

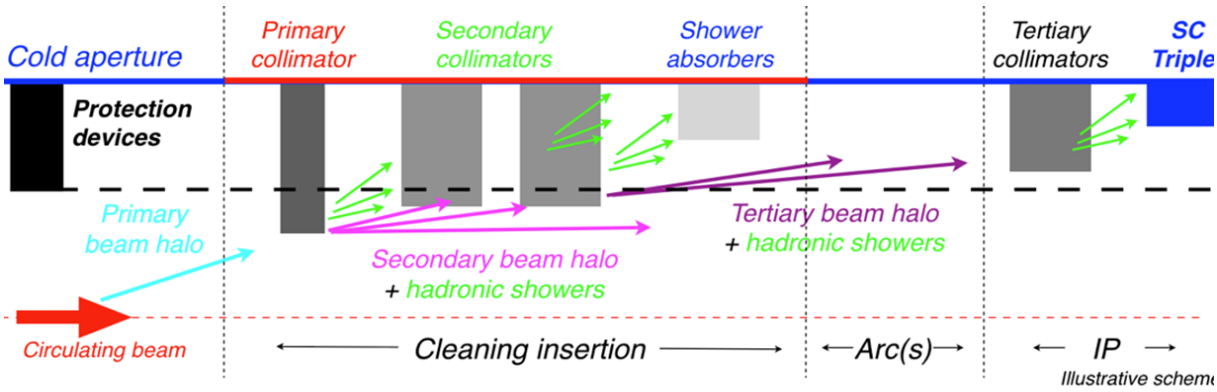

Fig. 8: Principle of multistage beam collimation (cleaning) at the LHC

Depending on their location in the ring and on their role, the collimator jaws are made of fibre-reinforced carbon, strongly absorbing copper and tungsten blocks. The system worked perfectly so far, also thanks to excellent beam stability and machine reproducibility. Around 99.99% of the protons that were lost from the beam were intercepted. A single set-up of the collimation system was required per year. From the machine protection perspective, the LHC collimator must fulfil a dual role, halo collimation (beam cleaning) and passive protection of the accelerator component.

The LHC beam dumping system (LBDS) is a complex system composed of 15 fast kicker magnets that deflect the beam to the outside of the ring, 15 normal-conducting septum magnets that deflect the beam vertically out of the plane of the ring, 10 dilution kicker magnets that spread the beam over the beam dump surface and finally the beam dump block [1, 2, 12, 13]. The layout of the LBDS is shown in Fig. 9. The 10 m long carbon dump block is the only LHC element capable of absorbing the nominal beam. The beam is swept over the dump surface to lower the power density. Without the sweep the beam could drill a hole with a depth of a few metres into the dump block through hydro-dynamic tunnelling [22].

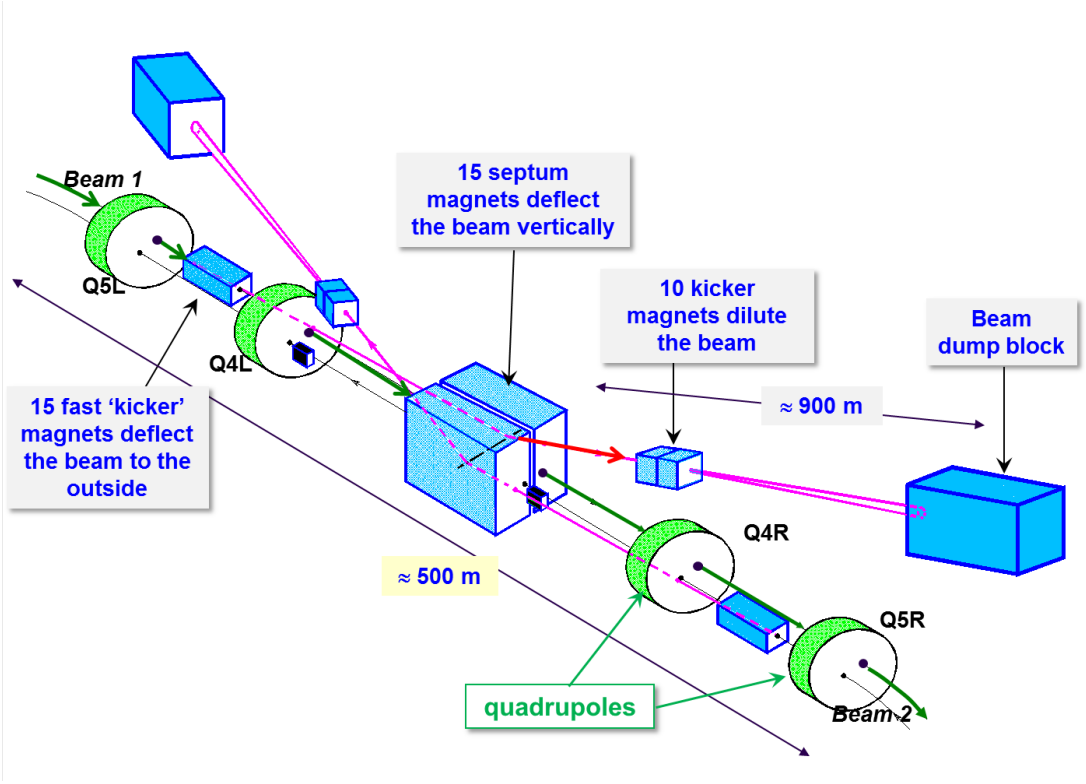

Fig. 9: Layout of the LHC beam dumping system (LBDS). For each beam 15 fast kicker magnets deflect the beam out of the ring into septum magnets that deflect the beam further into the vertical plane. On their way to the beam dump, fast dilution kickers spread out the bunches on the surface of the dump block (courtesy of M. Gyr).

A 3 μs long particle-free gap in the beam (beam abort gap) provides a time window for the LBDS dump kickers to raise to their nominal field. The dump kickers must be accurately synchronized to the beam abort gap to avoid spreading beam across the aperture during the kicker rise time. Possible failure modes are:

– the abort gap fills with beams (RF fault, debunching, injection error);
– the kicker synchronization fails or a kicker fires spontaneously (not synchronized) – so-called *asynchronous beam dumps*.

The asynchronous dump is the *ultimate unavoidable failure*: the LHC must be protected from this failure *passively* by absorbers. Two large absorbers in front of the extraction septum and in front of the first superconducting magnet downstream of the dump kickers and septa protect the LHC against damage from asynchronous dumps and from residual beam in the abort gap. Dump kicker powering, synchronization and triggering are designed to exclude out-of-sync triggers with high reliability. The spontaneous trigger is estimated to occur roughly once per year. So far none has been observed during high-intensity operation, but the system operated at reduced high voltage (4 TeV instead of 7 TeV), which also lowers the probability of spontaneous kicker firing.

A direct link between the LBDS and the injection system ensures that the LHC injection kicker magnets cannot be triggered with a timing that would inject beam into the abort gap. The beam population in the abort gap is monitored using synchrotron light. When the particle density in the abort gap exceeds a pre-defined threshold, the transverse feedback system of the LHC is used to excite those particles until they hit the collimators.

3 Interlock masks for commissioning and operation

Already during the design phase of the LHC MPS, the need for masking interlocks was recognized [12, 13]. Some flexibility is always required for low-intensity commissioning and setting up phases where the risk of damage is also much lower. To avoid masking interlocks by raising thresholds and opening tolerance windows for many parameters (with the risk of errors during the reversal), the concept of masking interlocks when *the beam is safe* was introduced at an early stage. Such a safe beam should not be able to damage accelerator components. The corresponding beam intensity limit depends on the beam energy and on the beam emittance. It also depends on the material that is considered for damage. At the LHC, copper was selected as reference material. Given the large difference between the energy required to quench a superconducting magnet and the energy required to damage the same magnet, it is clear that even a safe beam will be able to quench one or more magnets.

In the early design phase of the LHC machine protection system, around the years 2000–2004, not much was known about equipment damage by high-energy and high-intensity beams. In order to benchmark damage by a high-energy LHC-type beam, a dedicated experiment was set up in a SPS transfer line [23]. In this controlled experiment beams of 450 GeV protons with LHC bunch structure (25 ns spacing) and nominal LHC beam emittances were directed into a special target composed of a sandwich of materials (copper, stainless steel and zinc). From this experiment the damage limit for copper was established at 2×10^{12} protons. It should be noted that the damage limit for stainless steel was more than a factor of four higher. The observations were in good agreement with rather simple estimates for the damage limit based on shower simulations to define the deposited energy and heating of the materials up to the melting point.

FLUKA simulations [24] were used to extrapolate the experimental results from 450 GeV to the LHC operating energy of 7 TeV. The simulations predicted the following scaling law for the safe beam intensity I_{SB} with beam energy E:

- larger energy deposition implies a scaling $\propto 1/E$;
- shrinking of the transverse emittance (beam area) implies a scaling $\propto 1/E$;
- increase of the shower length provides some dilution $\propto \log(E)$.

The following effective scaling law with energy (at fixed normalized emittance) was obtained:

$$I_{\mathrm{SB}} \propto E^{-1.7}. \tag{2}$$

This equation is implemented inside the LHC safe machine system (SMP) and is indicated as the *Normal* set-up beam flag (SBF) in Fig. 10. At injection the intensity limit was set initially to 10^{12} charges, but was later lowered to 5×10^{11} because the beam emittances were roughly a factor of two smaller than expected in the LHC design. The SMP system is connected to reliable Beam Current Transformers (BCTs) and energy sources (based on the dipole fields with four-fold redundancy). It generates the SBF that is distributed over the LHC machine timing system to the BIS. The SBF states are:

- true = set-up beam: maskable BIS input signals can be masked;
- false = unsafe beam: no BIS input signals may be masked.

The beam interlock system is configured to allow masking certain classes of interlocks (maskable) when the SBF is true.

During initial operation it was realized that the definition of the SBF was too restrictive for beam commissioning that required higher bunch charges than expected (nominal bunches with 10^{11} protons) due to systematic effects in the LHC beam instrumentation, mainly the beam position monitoring system. As a consequence, relaxed SBF equations were introduced to be able to mask certain interlocks with 2–3 nominal bunches at top energy as shown in Fig. 10. During commissioning periods a MPS expert can relax the equation of the SBF within the SMP system for certain tests.

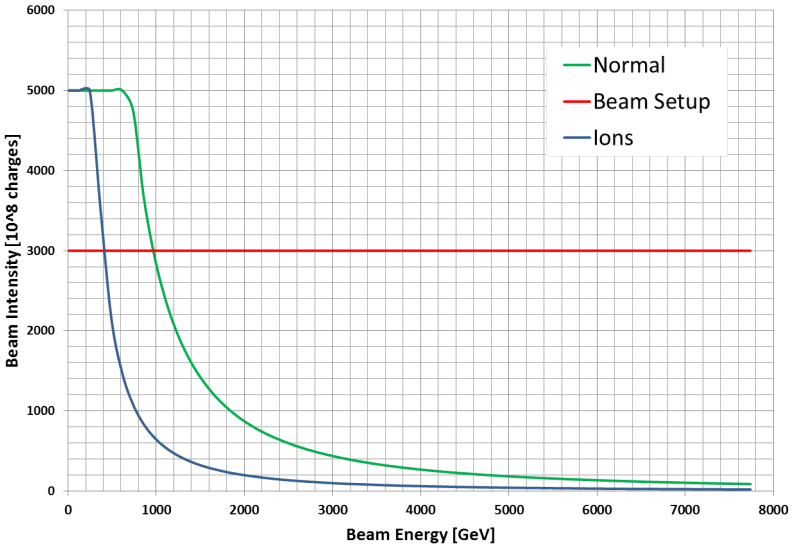

Fig. 10: Limiting beam intensity (in 10^8 charges) as a function of energy for the set-up beam flag. The *Normal* SBF corresponds to the scaling of Eq. (2). The *beam set-up* equation is a relaxed version used for beam commissioning and MPS set-up. For ion beams the limits are scaled by the ratio Z/A of the charge per nucleon.

4 Commissioning and organization

The MPS activities of the LHC were organized since the year 2000 inside a Machine Protection Working Group (later renamed to Machine Protection Panel, MPP). This group steered the design and followed up on implementation, issues and performance of the LHC MPS. All groups involved in MP activities were represented in the MPP.

An executive MP body, the restricted MPP (rMPP), was created when beam commissioning started. Each core MP system has one representative within the rMPP. The rMPP takes decisions related to MPS (example: BLM threshold changes) and steers the intensity ramp up of the machine. Its recommendations are submitted to the CERN management; in almost all cases the recommendations were accepted.

Before the machine startup, procedures were developed for the commissioning of the machine protection subsystems. The procedures contain test descriptions and frequency of tests (after stop or intervention). The procedures were then translated into a series of individual tests to be performed on the machine with and without beam, if required also at different intensity steps.

Machine protection tests are currently documented and tracked on a web page. One MPS expert of the commissioning team checks that all tests required for a given commissioning phase have been performed by the experts. It is in general the system expert that executes the tests for his system and not an independent person. This simple mechanism for tracking the commissioning will be improved in the future. The new concept with test tracking and electronic expert signatures is already in place for the commissioning of the LHC electrical circuits and magnets. The *AccTesting* framework [25] is based on pre-defined and agreed test sequences. Tests that are ready for execution can be scheduled for execution; test sequences are blocked until tests are analysed and signed by experts. The results are tracked and stored; no test step can be forgotten. Unfortunately due to lack of resources the test tracking was not yet implemented for the LHC MPS, but it is foreseen in the near future.

Passive protection of the LHC aperture with collimators and absorbers is a key ingredient for operating the LHC safely at high intensity. All failures affecting the machine on a global scale (orbit, optics, emittance, perturbations) must be intercepted by a protection device. The LHC machine setting

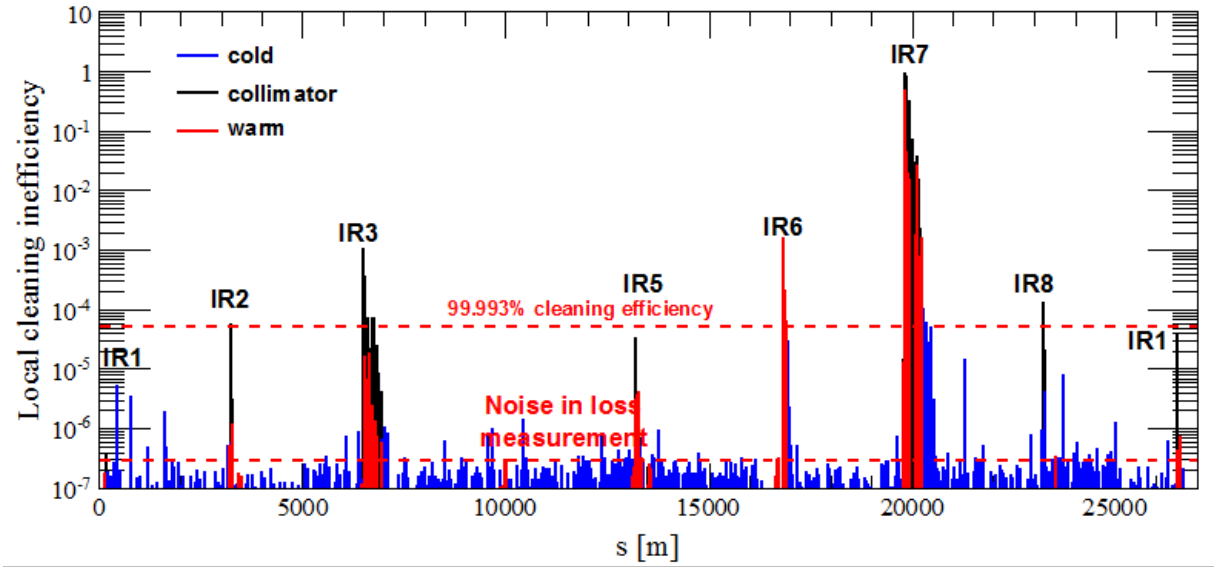

Fig. 11: Example of the beam loss distribution during a loss map performed at a beam energy of 4 TeV. The vertical scale gives the losses normalized to the peak loss at the primary collimator. The beam travels from left to right. Blue markers correspond to BLMs located on superconducting elements, red markers to room-temperature elements and black markers to collimators.

up involves:

- a well-corrected orbit (rms below 0.5 mm);
- a well-corrected optics (betatron function beating below 20% at injection and below 10% with colliding beams);
- a good knowledge of the aperture bottlenecks (after orbit and optics correction).

The machine aperture is measured after orbit and optics corrections have been finalized. The global aperture is measured at injection energy (typical aperture limit is at $\approx 12\ \sigma$, where σ is the beam size with a nominal normalized emittance of 3.5 μm). Local apertures are measured at injection (injection region, beam dump section and around the four experiments) and at top energy with fully squeezed beams (only the four experiments). At top energy the aperture depends on the optics at the different collision points. The aperture limit is generally located in the low-beta quadrupoles next to the collision points.

During machine set-up, all collimators and absorbers are aligned around the closed orbit with appropriate retractions. For good performance the orbit must be reproducible at the level of 50 μm. The machine set-up (orbit, optics, aperture and protection devices) is then validated by a campaign of loss maps and simulated asynchronous beam dump tests.

- During a *loss map* the transverse beam emittance is blown up until losses are observed on the collimators and absorbers; an example is shown in Fig. 11. The loss distribution provides a validation of the collimator alignment and hierarchy [20]. Initially the emittance was blown up by crossing the third-order resonance in the horizontal or vertical plane. This technique was however not always very reproducible; sometimes the losses were too small, sometimes they were massive and a large fraction of the beam was lost. From 2012 onwards emittance blow up was obtained by noise excitation using the transverse feedback system. This provided fine control over the losses and blow up could be applied to individual bunches [21].
- For a simulated *asynchronous test dump* a low-intensity beam (typically 1–3 bunches) is debunched by switching of the RF system. After a few minutes the particles drift into the abort gap,

the density of particles in the gap being monitored by a dedicated device based on synchrotron light. When the population in the abort gap is sufficient, a dump is triggered. The beam present in the region of the abort gap mimics the effect of an asynchronous dump. The loss distribution along the ring provides a validation of the dump protection alignment.

5 Intensity ramp

The entire LHC cycle, from injection to collisions, is always set up with low-intensity beams. The maximum intensity is three bunches with nominal intensity ($\approx 1.1 \times 10^{11}$ protons). The set-up is made with less than 1 permill of the nominal intensity, which represents a challenge for the beam instrumentation, since there should not be a significant bias of measurements between the set-up and the nominal full-intensity beam (for example for the beam position monitoring).

The intensity increase is steered through the restricted Machine Protection Panel (MPPr), which defines the intensity steps and the requirements to proceed with the intensity increase. The plan for the first learning year in 2010 foresaw three phases:

- low-intensity beams for commissioning and early experience. This phase was followed by an internal review of the MPS performance;
- an intensity ramp up to a stored energy of 1–2 MJ followed by 4 weeks of operation at that stored energy. Such a stored energy corresponded to state of the art in 2010. This phase was followed by an external review of the MPS performance;
- increase of the stored energy into the regime above 10 MJ.

Figure 12 shows the evolution of the peak stored energy for one LHC beam between 2010 and 2012. The slow ramp up of 2010 is clearly visible. With the experience that was accumulated in 2010, the stored energy was ramped up above 100 MJ in 2011 over 11 intensity steps, up to a maximum of 1300 bunches. The 2011 intensity ramp up took around 9 effective weeks; the rate was dictated by operational (and not MP) issues as soon as around 600 bunches were stored [4]. Losses and the need to adjust of BLM thresholds, vacuum element heating by the beam, beam stability, etc slowed down the pace. Finally, the 2012 intensity ramp up took just 2 weeks with seven intensity steps, as can be seen in Fig. 13.

6 Beam losses at the LHC

6.1 Witness beam

Injection into a synchrotron that has no circulating beam has the same MP issues as a linac [17]. Injection of an intense beam can represent a serious risk or require very important monitoring efforts (all power converters etc) to ensure that the beam is not lost directly on the aperture during the first turn, before the MPS is able to react. To overcome this issue, the concept of *witness* beam (or *beam presence*) was introduced for the LHC [26]. The principle is that only a probe bunch (typically 10^{10} protons, max 10^{11} protons) may be injected into an *empty* ring. Such a low intensity cannot provoke damage to components. A high-intensity injection requires a minimum beam intensity to be circulating; this is the best guarantee that the conditions are reasonable to avoid failures happening on the first turn, before the MPS is able to react. At the LHC the concept is based on a highly reliable and redundant intensity measurement. A flag indicating the beam presence (true/false) is transmitted to the extraction interlock system of the SPS injector where it is combined with a flag indicating that the SPS beam is has a probe intensity (max 10^{11} protons).

Despite storing up to 140 MJ at 4 TeV, not a single superconducting magnet was quenched at the LHC with circulating beam, despite quench thresholds of only a few tens of mJ. On the other hand, many magnets were quenched at injection, mainly due to (expected) injection kicker failures (seven events in 2012). The beam (roughly 2 MJ) is safely absorbed in injection dump blocks, but the shower leakage can quench magnets over a distance of around 1 km.

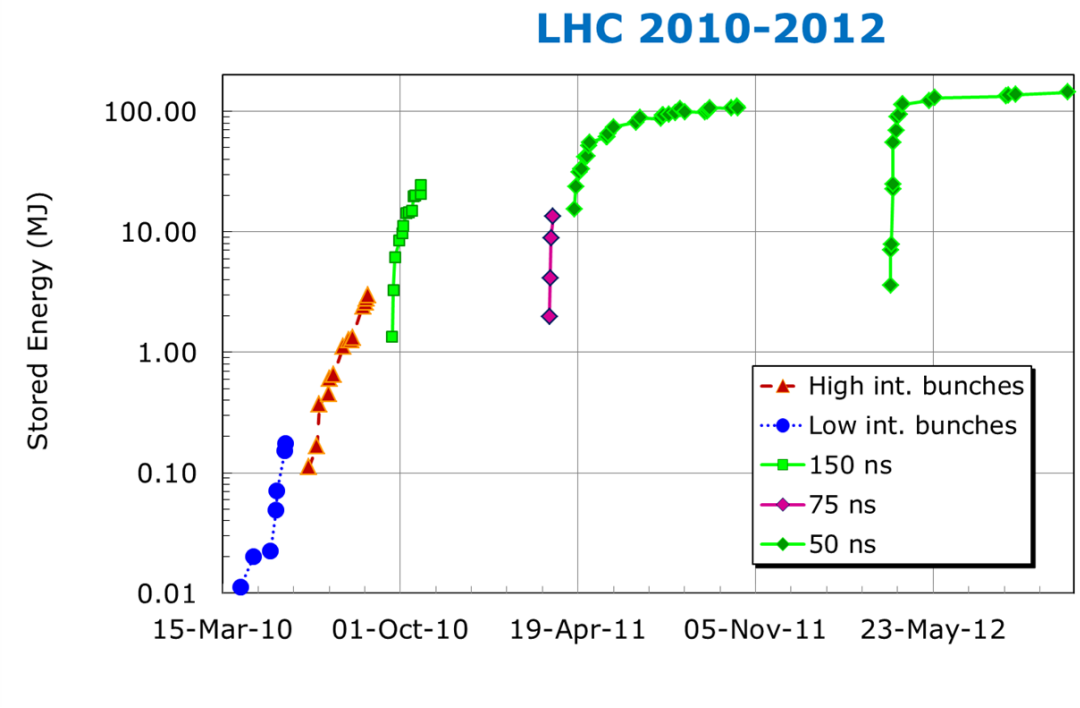

Fig. 12: Evolution of the peak stored energy in one LHC beam during Run 1. The different colours detail the beam structure in single bunches (first two periods) or in trains with bunch spacings of 150, 75 and 50 ns.

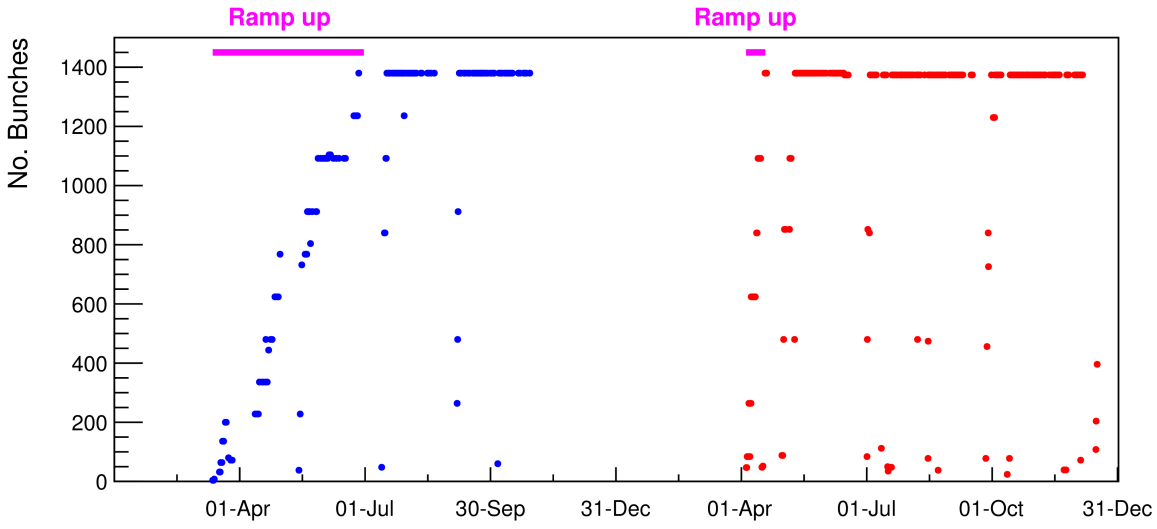

Fig. 13: Intensity ramp up of the LHC beams, expressed here through the number of bunches in 2011 and 2012

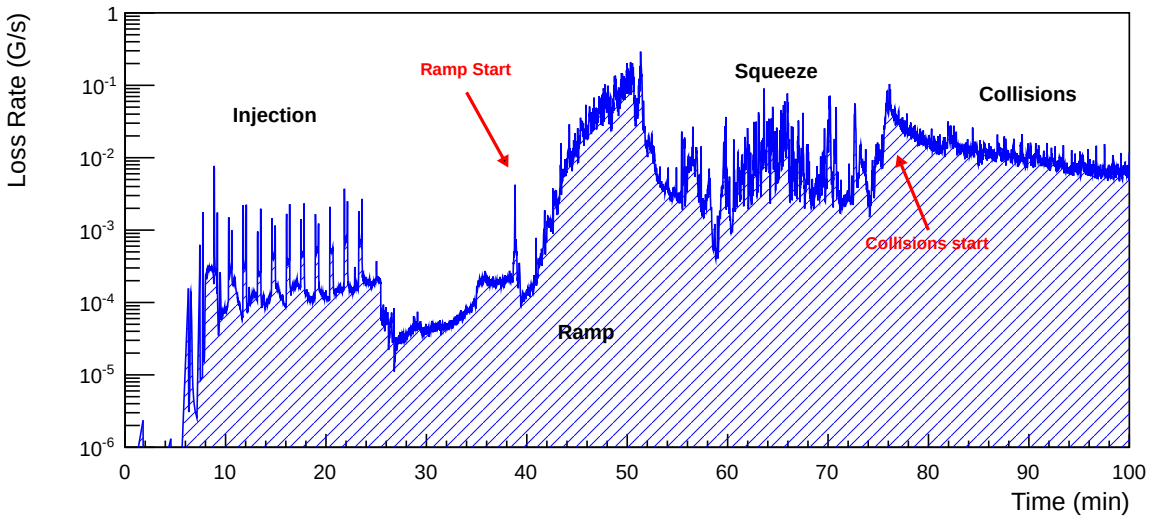

Fig. 14: Beam loss rates averaged over 1 s at the primary collimator along a typical LHC cycle in 2012. During the squeeze phase the optics around the experiments is changed to reduce (squeeze) the betatron function at the IP.

6.2 Beam loss in the cycle

At the LHC characteristic beam losses are observed in the various phases of the machine cycle. Those losses are part of regular operation and must be tolerated, even if one tries to minimize them. The normal losses are:

- injection losses (tails of the injected beams, injection oscillations, de-bunched beam);
- start of ramp losses (uncaptured beam loss);
- scraping on collimators (gap changes, orbit and tune shifts);
- losses from the beam halo when beams start to collide (beam–beam effect);
- losses due to the beam burn-off that are proportional to luminosity and performance.

The different types of losses are shown in Fig. 14 for a typical LHC cycle. The importance of the primary collimator opening appears clearly if the total beam intensity transmission presented in Fig. 15 is compared between 2011 and 2012. In 2012 the collimators were set closer to the beam in order to protect a smaller aperture, which allowed smaller betatron functions to be reached at the IP and therefore smaller beam sizes at the collision points, yielding a 60% higher luminosity [27]. The gap width has a strong impact on beam transmission and losses in the cycle.

- In 2011 the intensity losses during the ramp and squeeze phases are so small that they cannot be measured with the current transformers. The losses are completely dominated by collisions. The primary collimator gaps correspond to an opening of $\pm 7.5\ \sigma$, where σ is the real beam size.
- In 2012, on the other hand, there are already significant losses during the ramp when the collimator gaps are closed to their high-energy setting. There are noticeable losses in the optics squeeze due to an increased sensitivity to even small orbit jitter at the level of 50 μm. The primary collimator gaps correspond to an opening of $\pm 5.2\ \sigma$, where σ is the real beam size.

The regular loss distribution during collisions is shown in Fig. 16. The beam loss rates near the experiments are almost as high as at the collimators. This is due to the fact that around the experiments the loss monitors record collision debris, the physics processes at very small angles that are not covered by the experiments. Their signals are proportional to the collider luminosity.

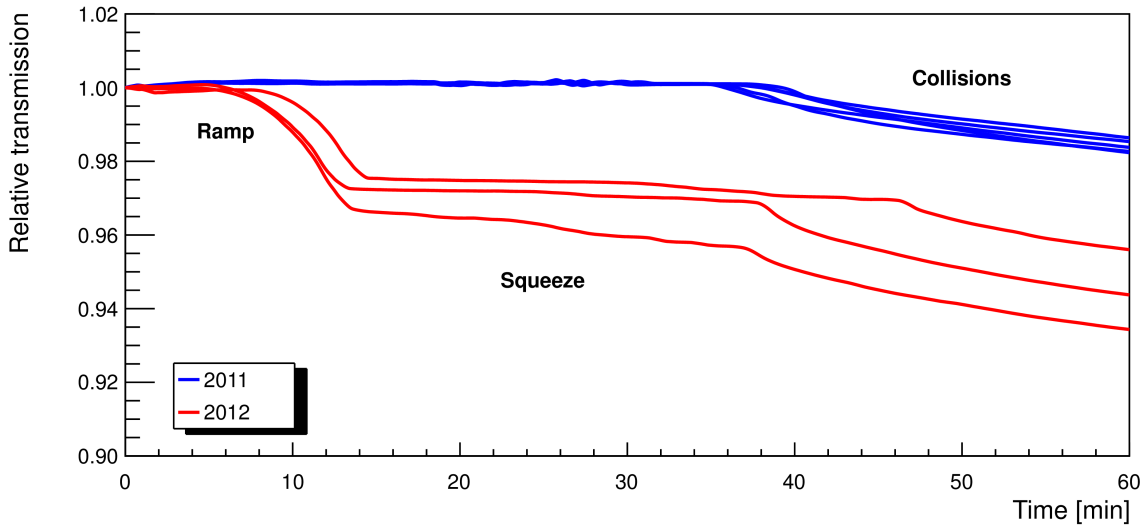

Fig. 15: Intensity transmission through the LHC cycle including the energy ramp and the betatron squeeze, the operation phase when the IP betatron functions are reduced to their value in collision.

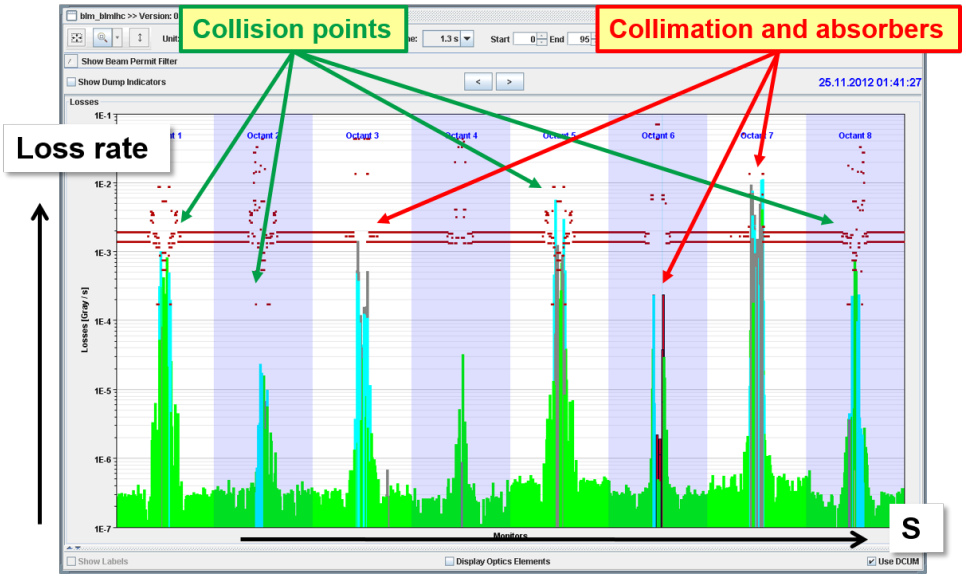

Fig. 16: Distribution of steady-state beam loss around the LHC ring with stable colliding beams

The thresholds of the LHC loss monitors are set to prevent quenches for monitors installed on superconducting elements and to prevent damage for room-temperature elements (for example collimators) [19]. In both cases some safety margin is desired. The initial thresholds were set before LHC operation started based on rather coarse quench level estimates coupled to GEANT, FLUKA and MARS simulations [24]. During the first years of operation the thresholds were progressively adapted (many were increased) based on experience. Initially the thresholds on collimators were set to limit the average power loss significantly below the peak design power of 500 kW. The thresholds were only increased to match the nominal power loss once the operation of the LHC was well controlled and understood. Quench tests with wire scanners (nice point-like particle source), orbit bumps and short and high losses in the collimation area were used to determine more accurately the quench limits [19].

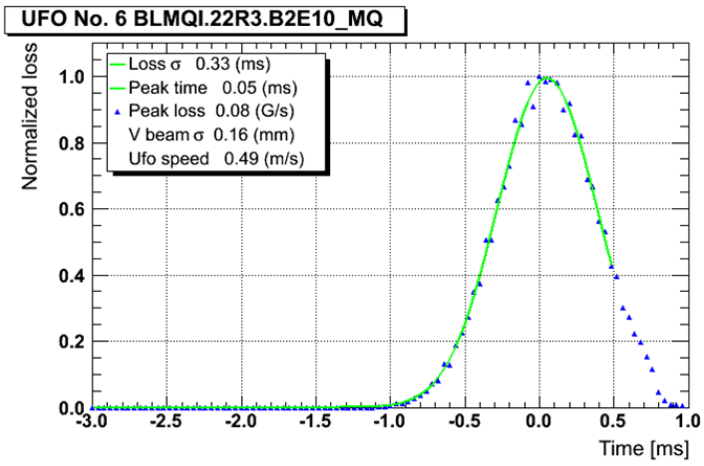

Fig. 17: Time evolution of the beam loss during a UFO event. The vertical scale represents the beam loss relative to the highest loss during the UFO event.

6.3 UFO losses

Very fast beam loss events (time-scale of a millisecond and below; see Fig. 17) mainly in superconducting regions have been the *surprise of LHC operation*. Those fast losses were nicknamed UFOs (unidentified falling objects) when it became clear that the UFO-type beam losses were due to small objects falling into the beam, the subsequent interaction of the beam particles leading to the showers and the beam losses [28, 29]. The BLM signals are consistent with small (tens of μm diameter) dust particles 'entering' the beam. The vast majority of UFO events lead to losses below dump threshold, but around 20 beam dumps were triggered by such UFO-type events every year between 2010 and 2012. UFO events localized within the injection kickers could be traced to aluminium oxide dust present in large quantities on the inner surface of the ceramic vacuum chamber of those elements. A cleaning campaign of the kickers was made during the long shutdown in 2013–2014. UFOs occur mainly with high-intensity beams, but there is conditioning with beam; the event rate drops from 10 to 2 per hour over a year as shown in Fig. 18. To improve the sensitivity of the LHC arc BLMs to UFO events, two out of six BLMs installed around the arc quadrupoles were re-located to the dipoles during the shutdown [30]. The monitoring and protection (quench prevention) capabilities of the BLMs are significantly improved with this re-location for the coming 6.5 TeV run. While the beam losses will increase due to the higher beam energy, the quench thresholds of the magnets will come down by a factor of 3–4, making the situation at 6.5 TeV more critical.

7 Diagnostics and control

The three main requirements for a modern MPS are:

- protect the machine: the highest priority is obviously to avoid damage to the accelerator;
- protect the beam: complex protection systems reduce the availability of the accelerator; the number of 'false' interlocks stopping operation must be minimized. This implies sometimes a trade-off between protection and operation;
- provide the evidence: clear (post-mortem) diagnostics must be provided when the protection systems stop operation or when something goes wrong (failure, damage, but also 'near misses').

Once the MPS components have been commissioned, it is essential that the protection functionality is maintained during operation. Automated checks of the MPS components as pre- or post-flight checks

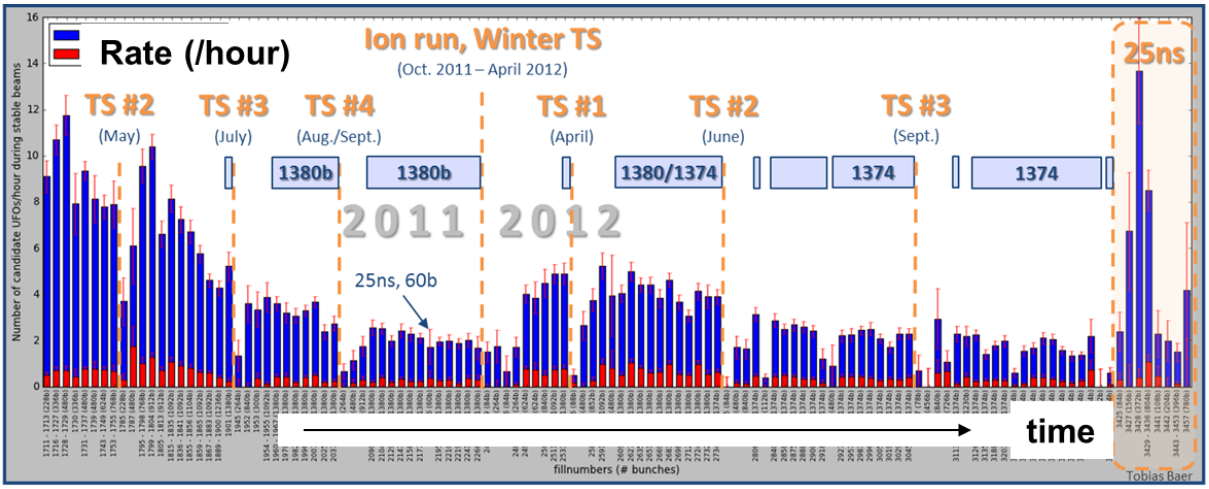

Fig. 18: Evolution of the hourly rate of UFO events in 2011 and 2012. A slow conditioning is observed during each run (adapted from [28, 29]).

can ensure that the MPS functionality is not degraded. For colliders with long cycle times there are two types of checks that fit well into the cycle, namely pre-flight checks before injection and post-flight checks on data collected during a fill or during the beam dump (post-mortem data). Such tests can come in multiple forms, for example the verification of MPS-related settings such as interlock thresholds, configuration checks of the beam interlock systems, automated analysis of the faults and MPS reaction chain and automated analysis of the dump system action. At the LHC the BLM system integrity checks must be performed at least once per 24 h period (or after the following beam dump if that interval is exceeded) [31]. The integrity checks are performed by a high-voltage modulation that is analysed to detect 'dead' channels.

At the LHC the MPS is so critical that for every beam dump post-operation checks (POCs) are performed on the beam dump system post-mortem data (equipment and beam signals). Data collection and analysis are triggered automatically after each dump. The analysis covers internal beam dump system signals and external beam information like dumped beam intensities, losses and beam positions in and around the dump channel. The analysis ensures that all signals are correct and that there is no loss of redundancy; the beam dumping system can then be considered 'as good as new'. Machine operation is stopped if the beam dumping system POCs fail; an expert verification is required to re-start beam operation. This concept was so successful that it was extended to LHC injection in the form of automated checks of each injection quality.

Already during the design of the LHC MPS, post-mortem (PM) diagnostics was identified as a key component to understand the root causes of beam dumps [32]. All key LHC systems implement circular PM data buffers that are frozen and read out when a beam dump is triggered. The PM trigger is transmitted over the LHC machine timing system. The PM system also inserts an automatic entry in the LHC electronic logbook. Sampling frequencies of the data in the PM buffers range from ms or turn level to tens of milliseconds and are adapted to each system. Synchronization is critical to make sense of the data and define the event sequence; therefore time stamping is made at the source of the data. The LHC post-mortem data volume is currently around 200 MBytes. Since the analysis of the large volume of LHC PM data can be tedious, automatic analysis tools have been implemented to support operators and MPS experts in their work. The LHC PM server and graphical user interfaces are based on the JAVA language with standard interfaces to extract raw data and provide analysed data [33]. This allows many persons to contribute to the analysis modules. After a beam dump, injection into the LHC is blocked until the PM data is collected, pre-analysed (automatic) and signed by the operation crews. If an automated

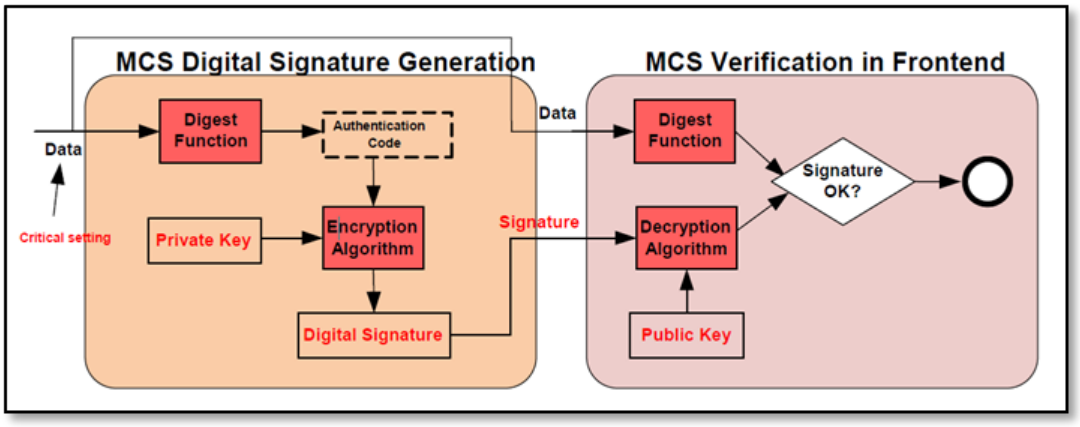

Fig. 19: Principle of management of critical settings (MCS) with digital signature of the data

analysis identifies a critical problem, injection can only be released by a MPS expert. To detect any possible long-term trend MPS experts re-analyse all PM events collected above injection energy within a few days. This also provides a view that is independent of the operation teams.

7.1 Settings control

Depending on its size, complexity and energy range an accelerator has a large volume of MPS settings in the form of BLM thresholds, current references and tolerances, etc. The management of settings associated to the MPS poses special problems, since access to the settings must be limited to authorized experts, yet it must be possible to inspect and download the settings from a central database to the front-end systems of the MPS. For the LHC access restrictions have been put in place based on *role-based access control* (RBAC) [34] and on digital signatures. To protect MPS settings the concept of *management of critical settings* (MCS) [35] was developed, a settings protection system that is fully embedded in the controls middleware and settings management. A setting that is defined as *critical* has an associated digital signature. Only a user that has the appropriate RBAC role (MPS expert, BLM expert, etc) is able to generate the digital signature. The digital signature is generated at the moment when a setting is changed by an authenticated expert. The setting and the associated digital signature are transmitted together to the front-end computer: a critical setting is only accepted with a valid digital signature, see Fig. 19.

7.2 Software interlocks

The LHC MPS has inputs from many systems that operate independently and there is only very limited information exchange between the various systems. To implement interlocks on a global machine scale with correlation of data between many LHC systems, a software interlock system (SIS) was developed [36, 37]. This system is able to collect information from any control device of the LHC and its injectors. The LHC SIS can dump one or both beams and inhibit LHC injection. It may perform global scale analysis among systems, correlate injector and LHC ring data for injection protection, etc. New interlock tests can be implemented rapidly to protect against unexpected issues. The LHC SIS is by design rather slow (reaction time of the level of seconds) but it is able to detect anomalies that could lead to problems in the future or prevent unnecessary beam dumps at injection. The SIS is based on a JAVA core server, with a large data buffer for currently up to around 2500 devices and settings. For each device data a timeout and no-data policy is defined. Over 5000 tests are defined with the LHC SIS; tests can be defined as simple value comparisons or complex JAVA logic.

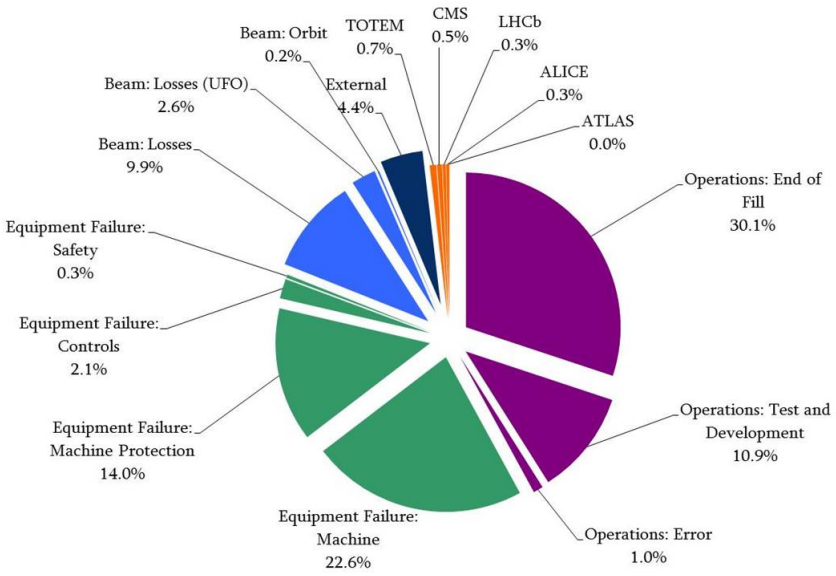

Fig. 20: Distribution of beam dump causes during the 2012 LHC proton run (adapted from [38]). The statistics include only beam dump above injection energy.

8 Machine availability

Peak performance in terms of luminosity is not the most important LHC parameter; the integrated performance is much more critical. To optimize the integrated time with collisions it is important to maximize the time spent colliding the beams and to minimize the time to re-establish collisions at the end of a fill (turn-around time). The fraction of time that the LHC spends colliding beams stably for the experiments amounted to 30–35% of the total scheduled operation time [4]. The average duration of collisions was only 6 h while the optimum duration would be in the range of 8 to 12 h. The reason why the fills are so short is because during high-intensity operation only one out of three fills is dumped by the operation crews; see Fig. 20. The other two-thirds of the fills are dumped by the LHC MPS [38]. Roughly 14% of the beam dumps are due to the failures of MPS subsystems ('false' dumps), with the following distribution of the MPS subsystems:

- quench protection system (radiation to electronics): 65%;
- BLM system: 13%;
- beam dumping system: 12%;
- software interlock system: 5%;
- powering interlock system: 2.5%;
- beam interlock system: 1.5%.

A reliability working group predicted the rate of false dumps and the safety of the LHC MPS for 7 TeV operation before the LHC was switched on. This can now be compared with observations, but it must be observed that so far LHC operated at 4 TeV while the predictions were made for 7 TeV, which is not completely equivalent for certain systems. The observations are roughly in line with predictions, but some failures do not match completely; in particular, radiation to electronics was not included in the initial predictions [39, 40].

9 Machine experiments

Machine experiments can be very exciting, but also risky periods for an accelerator. Frequently the machine is operated at some 'distance' from standard conditions. For example collimator settings, orbit and optics may be changed. At the LHC every experiment is categorized according to the foreseen changes to the machine and to the beam intensity. Experiments using intensities above the SBF limit must be prepared with a detailed written description of the changes to machines and the test procedures. In many cases the analysis of the document helped improve the efficiency of the experiment by spotting 'impossible' things. This encourages experimenters to think about options with smaller MP footprint, for example lower intensity.

10 Summary

The LHC MPS has been very successful in protecting the LHC with over 100 MJ of stored beam energy while the LHC was operated at 4 TeV. No component was damaged by a failure leading to beam loss; the MPS therefore fulfilled its job. As expected, operation of the LHC was significantly constrained by MP due to the high stored energy and very low quench levels. Nevertheless, operation was rather smooth and the intensity ramp up from well below 1 MJ during the commissioning phase to over 100 MJ stored energy took only 2 weeks in 2012. In 2015 the beam energy will be increased to 6.5 TeV: the stored energy of the beams will increase by a factor of 2 to 3, while the quench levels of the magnet drop by a factor of 3 to 5. This will be a new challenge for LHC operation, while at the same time the focus is shifting more and more towards high(er) machine availability.

References

[1] L. Evans, *The Large Hadron Collider: A Marvel of Technology* (EPFL Press, Lausanne, 2009).

[2] O. Brüning et al. (eds), LHC Design Report, CERN-2004-003 (2004). http://dx.doi.org/10.5170/CERN-2004-003-V-1.

[3] H. Pfeffer, Protection of hardware: powering systems (PC, NC and SC magnets), these proceedings.

[4] J. Wenninger et al., Operation and configuration of the LHC in Run 1, CERN-ACC-NOTE-2013-0041 (2013).

[5] L. Rossi, The Large Hadron Collider of CERN and the roadmap toward higher performance, CERN-ACC-2014-0225.

[6] R. Schmidt et al., Machine protection challenges for HL-LHC, Proc. 5th International Particle Accelerator Conference, IPAC14, Dresden, Germany, `http://www.jacow.org/`.

[7] F. Zimmermann, M. Benedikt, D. Schulte and J. Wenninger, Challenges for highest energy circular colliders, Proc. 5th International Particle Accelerator Conference, IPAC14, Dresden, Germany, `http://www.jacow.org/`;
Future circular collider study, https://espace2013.cern.ch/fcc/Pages/default.aspx.

[8] F. Bordry, Status of the Large Hadron Collider (LHC), Proc. 2008 European Particle Accelerator Conference, EPAC08, Genoa, Italy, 2008, `http://www.jacow.org/`.

[9] J. Wenninger, Status of the LHC, Proc. 2009 Particle Accelerator Conference, PAC09, Vancouver, Canada, 2009, `http://www.jacow.org/`.

[10] F. Bordry et al., The first long shutdown (LS1) for the LHC, Proc. 4th International Particle Accelerator Conference, IPAC13, Shanghai, China, 2013, `http://www.jacow.org/`.

[11] J.Ph. Tock et al., Consolidation of the LHC superconducting circuits: a major step towards 14 TeV collisions, Proc. 4th International Particle Accelerator Conference, IPAC13, Shanghai, China, 2013, `http://www.jacow.org/`.

[12] R. Schmidt et al., *New J. Phys.* **8** (2006) 290. http://dx.doi.org/10.1088/1367-2630/8/11/290

[13] R. Schmidt et al., Machine protection and interlock systems for LHC, these proceedings.

[14] R.B. Appleby et al., *Phys. Rev. ST Accel. Beams* **13** (2010) 061002. http://dx.doi.org/10.1103/PhysRevSTAB.13.061002

[15] V. Kain, Beam dynamics and beam losses for circular machines, these proceedings.

[16] M. Werner et al., A fast magnet current change monitor for machine protection in HERA, Proc. ICALEPS 2005, Geneva, Switzerland, http://www.jacow.org/.

[17] V. Kain, Beam transfer and machine protection, these proceedings.

[18] B. Dehning, LHC beam loss monitor system, these proceedings.

[19] B. Auchmann et al., *Phys. Rev. ST Accel. Beams* **18** (2015) 061002. http://dx.doi.org/10.1103/PhysRevSTAB.18.061002

[20] S. Redaelli, Beam cleaning and collimation systems, these proceedings.

[21] B. Salvachua et al., Handling 1 MW losses with the LHC collimation system, Proc. 5th International Particle Accelerator Conference, IPAC14, Dresden, Germany, http://www.jacow.org/.

[22] A. Bertarelli, Beam induced damage mechanisms and their calculation, these proceedings.

[23] V. Kain et al., Material damage test with 450 GeV LHC-type beam, LHC-Project-Report-822 (2005).

[24] N. Mohkov and F. Cerutti, Beam material interaction, heating and activation, these proceedings.

[25] A.A. Gorzawski et al., The AccTesting framework: an extensible framework for accelerator commissioning and systematic testing, Proc. ICALEPS 2013, San Francisco, CA, USA, http://www.jacow.org/.

[26] R. Schmidt and J. Wenninger, LHC injection scenarios, LHC Project Note 287 (2002).

[27] S. Redaelli et al., Operation of the betatron squeeze at the LHC, Proc. 4th International Particle Accelerator Conference, IPAC13, Shanghai, China, http://www.jacow.org/.

[28] T. Baer, Very fast losses of the circulating LHC beam, their mitigation and machine protection, CERN-THESIS-2013-233.

[29] T. Baer et al., UFOs in the LHC: observations, studies and extrapolations, Proc. 3rd International Accelerator Conference, IPAC12, New Orleans, LA, USA, http://www.jacow.org/.

[30] B. Auchmann, BLM threshold strategy, LHC Performance Workshop, Chamonix, 2014, http://indico.cern.ch/event/315665/.

[31] J. Emery et al., LHC BLM single channel connectivity test using the standard installation, Proc. European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators DIPAC09, Basel, Switzerland, http://www.jacow.org/.

[32] E. Ciapala, F. Rodriguez Mateos, R. Schmidt and J. Wenninger, The LHC post-mortem system, LHC Project Note 303 (2002).

[33] R. Gorbonosov et al., Plug-in based analysis framework for LHC post-mortem analysis, Proc. International Conference on Accelerator and Large Experimental Physics Control Systems ICALEPS 2013, San Francisco, CA, USA, http://www.jacow.org/.

[34] P. Charrue et al., Role-based access control for the accelerator control system at CERN, Proc. International Conference on Accelerator and Large Experimental Physics Control Systems ICALEPS 2007, Knoxville, TN, USA, http://www.jacow.org/.

[35] W. Sliwinski et al., Management of critical machine settings for accelerators at CERN, Proc. International Conference on Accelerator and Large Experimental Physics Control Systems ICALEPCS 2009, Kobe, Japan, http://www.jacow.org/.

[36] J. Wozniak, V. Baggiolini, D. Garcia Quintas and J. Wenninger, Software interlock system, Proc. International Conference on Accelerator and Large Experimental Physics Control Systems ICALEPS 2007, Knoxville, TN, http://www.jacow.org/.

[37] L. Ponce, J. Wenninger and J. Wozniak, Operational experience with the LHC software interlock system, Proc. International Conference on Accelerator and Large Experimental Physics Control

Systems ICALEPS 2013, San Francisco, CA, USA, http://www.jacow.org/.

[38] B. Todd, Performance and availability of MPS 2008–2012, Proc. 2013 MPP Workshop, CERN-ACC-2014-0041.

[39] R. Filippini et al., Reliability analysis of the LHC beam dumping system taking into account the operational experience during LHC Run 1, Proc. International Conference on Accelerator and Large Experimental Physics Control Systems ICALEPS 2013, San Francisco, CA, USA, http://www.jacow.org/.

[40] J. Uythoven, Dependability calculations prior to 2008 and operational experience during first LHC run, https://indico.cern.ch/event/277684.

[41] M. Bajko et al., Report of the task force on the incident of 19th September 2008 at the LHC, CERN-LHC-PROJECT-Report-1168.

Appendix A: Powering incident in 2008

Since the LHC powering incident in 2008 was the most severe damage that ever happened to an accelerator, a brief description of it is appended here even though this incident happened *without beam* [9].

In the morning of 19 September 2008 the last commissioning step of the main dipole circuit (154 magnets) of sector 34 (between IP3 and IP4 in Fig. 1) was started, a ramp to 9.3 kA which corresponds to a beam energy of 5.5 TeV. During the ramp an electrical fault developed in the powering busbar interconnection at a current of 8.7 kA. A resistive voltage appeared and increased to 1 V after less than 0.5 s, leading to the power converter trip. The current started to decrease in the circuit and the energy discharge switch opened, inserting dump resistors in the circuit. In this sequence of events, the quench detection, power converter and energy discharge systems behaved as expected. No resistive voltage appeared on the dipoles of the circuit, individually equipped with quench detectors and the quench of a magnet has been excluded as initial event.

Within the first second, a main electrical arc and multiple smaller secondary arcs developed and punctured the helium enclosure, leading to release of helium into the insulation vacuum of the cryostat. After a few seconds the beam vacuum also degraded. The spring-loaded relief discs on the insulation vacuum enclosure opened when the pressure exceeded atmospheric, thus relieving the helium to the tunnel. They were however unable to maintain the pressure rise below the nominal 0.15 MPa absolute, thus resulting in large pressure forces acting on the vacuum barriers separating neighbouring subsectors (a subsector corresponds to two 107 m long cells), which damaged them as can be seen in Fig. A.1. The forces displaced dipoles in the affected zone from their cold internal supports and knocked the short straight section (SSS) cryostats housing the quadrupoles and vacuum barriers from their external support jacks at three positions, in some locations breaking the anchors in the concrete floor of the tunnel. The displacement also damaged the connections to the cryogenic distribution line. The main damage zone extends over approximately 700 m.

About 2 tons of helium were rapidly released to the tunnel, producing a cloud which triggered oxygen deficiency hazard detectors and tripped an emergency stop, thus switching off all electrical power from sector 34. Before restoration of electrical power enabled the actuation of cryogenic valves, another 4 tons of helium were lost at lower flow rates. The total loss of inventory thus amounts to about 6 tons out of 15 tons initially in the sector.

A post-mortem analysis of cryogenic temperature data revealed a significant temperature anomaly in sector 34 during a powering step to 7 kA performed a few days before the incident. A steady temperature increase of up to 40 mK occurred in the cryogenic cell of the incident. The excess power in the incident cell corresponds to an unaccounted resistance of around $220\ \mathrm{n\Omega}$. Given the location of the primary electrical arc, the most likely hypothesis for the cause of the incident is a problem of the busbar joint. The structure of such a joint is shown in Fig. A.2. The joints are brazed but not clamped and the nominal joint resistance is $0.35\ \mathrm{n\Omega}$. The incident could be reproduced in simulation assuming a bad electrical and thermal contact of the copper stabilizer at the joint due to lack of solder or poor-quality brazing [41].

A total of 53 magnets, 39 dipoles and 14 quadrupole SSSs had to be removed from the tunnel and brought to the surface for cleaning and repair. Most of them were replaced with spare magnets. All magnets have been thoroughly re-tested before re-installation in the tunnel.

Both arc beam vacuum chambers were contaminated by soot from electrical arcs and chips of multilayer insulation over roughly 80% of their length as can be seen in Fig. A.3. Contamination by chips of multilayer insulation has been found over long distances away from the position of the original incident. These chips are deposited mostly on the beam screen surface, from where they are removed by in situ cleaning.

Following the incident and a review of the magnet and busbar protection system, a Quench Protection System (QPS) upgrade was launched to protect all busbar joints of the arc main dipole and main

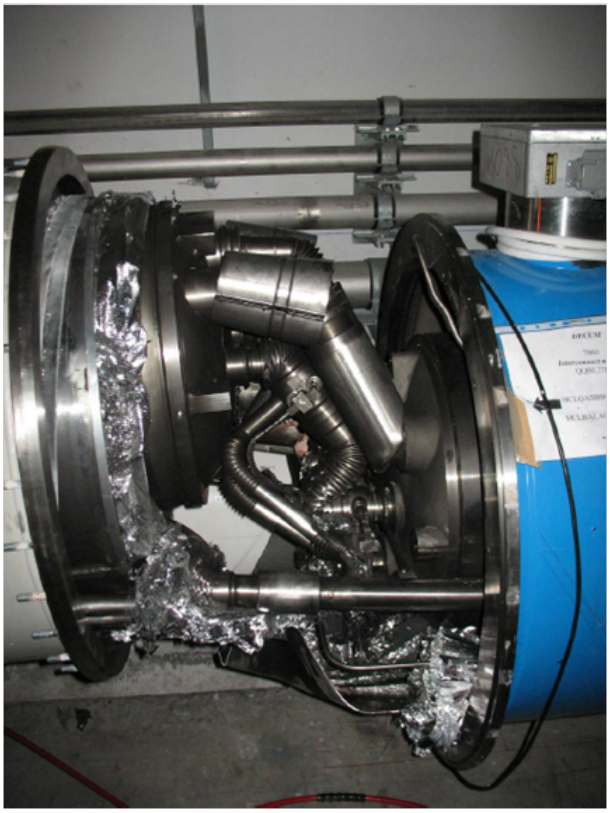

Fig. A.1: A damaged interconnect between a quadrupole and a dipole magnet

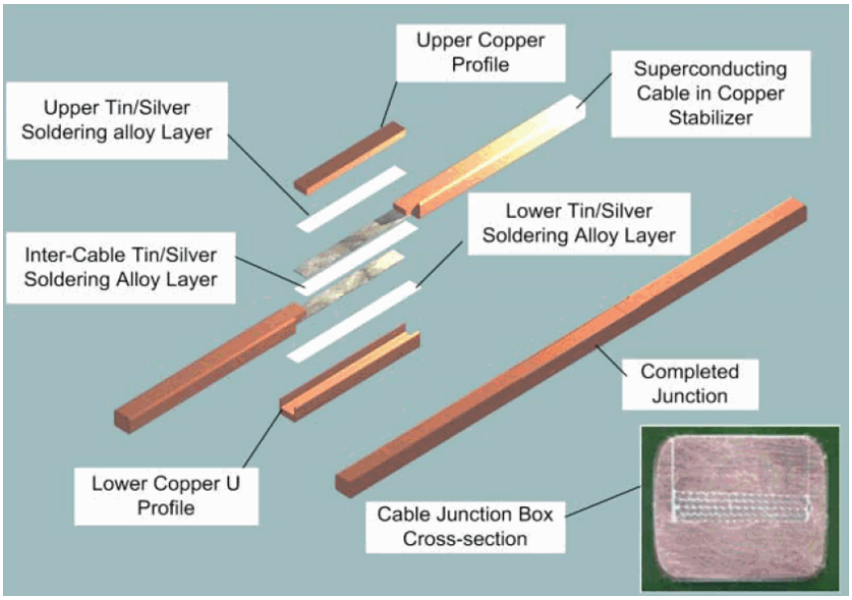
Upper Copper Profile
Superconducting Cable in Copper Stabilizer
Upper Tin/Silver Soldering alloy Layer
Lower Tin/Silver Soldering Alloy Layer
Inter-Cable Tin/Silver Soldering Alloy Layer
Completed Junction
Lower Copper U Profile
Cable Junction Box Cross-section

Fig. A.2: Schematic of the main dipole busbar joint. The superconducting cable is embedded in a copper stabilizer. At the joint the two busbar ends are inserted with solder into a copper profile and welded.

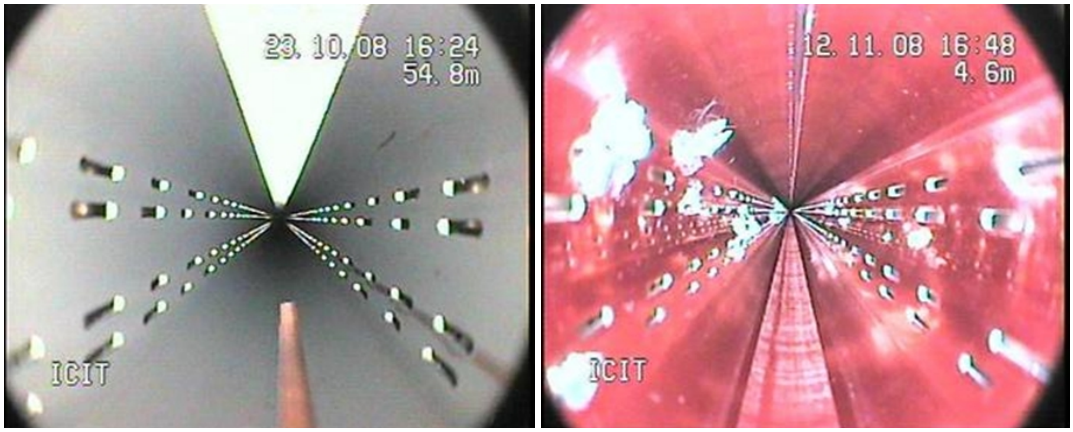

Fig. A.3: Example of a vacuum chamber beam screen covered with soot from the incident

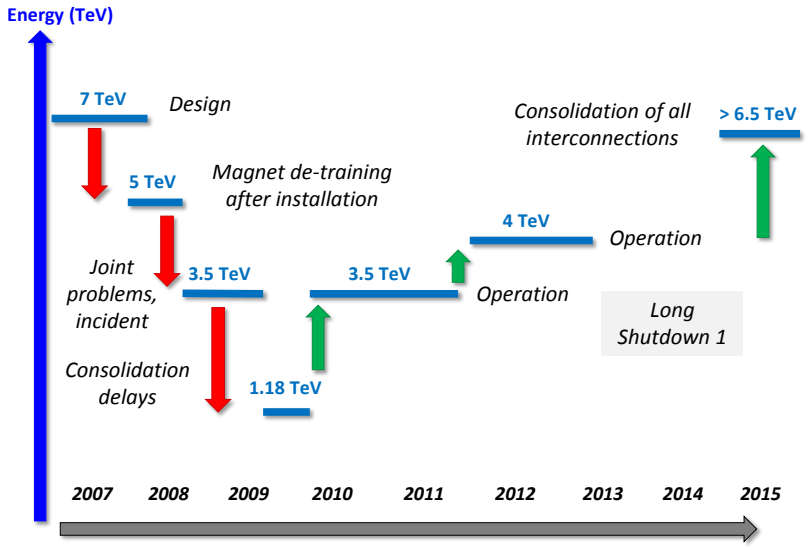

Fig. A.4: Evolution of the LHC beam energy from design to the present day

quadrupole circuits. The required voltage tabs were available, but a large volume of electronics had to be developed. The busbar protection was finally operational in the spring of 2010 when the LHC was commissioned at 3.5 TeV. The evolution of the LHC beam energy over time is shown in Fig. A.4, in 2015 the energy will be pushed to 6.5 TeV after the consolidation campaign of the 2 year long shutdown.